Particle Accelerators for High Energy Physics
A Short History

D. A. Edwards

April 5, 2006

Abstract

These four talks were prepared for the XXXIV International Meeting on Fundamental Particles entitled "From HERA and the Tevatron to the LHC" held at El Escorial, Spain, April 2–7 2006. The first - The Past - starts from the point some seventy-five years ago when it was recognized that earth-based particle energies beyond the capability of electric fields and radioactive sources were needed, and concludes with the remarkable achievement of proton-antiproton collisions. The second - Today - discusses the current multi-accelerator facilities, taking HERA and the Tevatron as examples of the advances in beam physics and, importantly, technology associated with their development. The third - Tomorrow - concentrates on the LHC and its challenges, and comments on its predecessor LEP in the same underground enclosure. Talk 4 - The Future - starts with steps toward an electron-positron linear collider design, and and then turns to some final remarks about the steps beyond.

1   The Past

The discoveries of the electron and proton were achieved through the use of gas discharge tubes and radioactive sources, by experiments conducted by the such illustrious names as J. J. Thompson and Lord Rutherford. But by the second decade of the twentieth century, it was clear that higher beam energies were needed to carry particle physics forward.

One direction was continuation of work on electrostatic fields, and that avenue lead to the Cockroft-Walton and van de Graaff developments. The single-pass character of these devices precludes particle energy above the 10 MeV or so range. The other path put forward at roughly the same time was the use of multiple-pass acceleration through the use of radiofrequency fields, and this is the direction that which will be discussed here and throughout these talks.

The half-century summarized in this first talk was the period during which virtually all the principles of accelerator science in present and near-future use in support of elementary particle physics were invented. In contrast, the progress outlined in the next two talks

required advances in technology both within and outside of the field to exploit the principles set forth in this earlier period.

As a *caveat* to the listener or reader, I should admit that my knowledge of the early years of this time is based on reading and conversations. I did not encounter a real accelerator until about 1948.

1.1 Introduction of Radiofrequency Fields

Repetitive energy increments by the same structure or identical structures require time-varying fields (if electromagnetic fields are to be used to provide the emf) according to Faraday's Law:

$$\oint \vec{E} \cdot \vec{ds} = - \int \dot{\vec{B}} \cdot \vec{dA} \tag{1}$$

where the line integral on the left extends through the repeat distance of the device, a single turn in the case of a circular accelerator. From the outset, it was clear that resonant systems provided a natural avenue for production of the accelerating fields, and so radiofrequency (RF) systems entered the accelerator world. The first operating example was constructed by Rolf Wideröe [1] in 1927 while a graduate student at the University of Aachen. This precursor of today's linacs is depicted in Fig. 1, which appeared in Wideröe's dissertation. A few megahertz was a high frequency in the 1920s. Although I could not find a statement of the frequency associated with the apparatus in Fig. 1, we can make an estimate. He was accelerating a mix of sodium and potassium ions to a total kinetic energy of 50 keV. So within the central drift tube, the energy was 25 keV, corresponding to a speed for sodium atoms of 0.6×10^6 m/sec. Wideröe says that the total length of the glass tube is 0.88 m, and the drift tube must be about 0.2 m long. For a half-period of the RF oscillation to elapse while the ion passes through the drift tube, the oscillator frequency must be about 1.5 MHz.

The invention of the cyclotron followed almost immediately. Ernest Lawrence decided that an indefinite extension of Wideröe's glass tube was unappealing, so the device built by Lawrence and (then graduate student) M. Stanley Livingston used a magnetic field to bring the particles repeatedly back through the same RF accelerating field. In January 1931, the first cyclotron produced 80 keV protons [2]. This device remains on exhibit at the Lawrence Hall of Science [3] and is shown in Fig. 2. The coin in the figure is a Nobel Prize medal. The oscillator frequency was in the neighborhood of 3 MHz.

The spiral orbits of a cyclotron were not attractive for scaling to high energies due to the increasingly massive magnets. Proton accelerators of this form did achieve 450 MeV or so, with frequency modulated RF systems to accommodate Lorentz factors that differed from unity; these were called synchrocyclotrons.

A return of Wideröe's linear concept was facilitated by the development of high frequency power sources for RADAR during World War II. A proton linear accelerator was constructed under the direction of Luis Alvarez at Berkeley using 200 MHz war-surplus
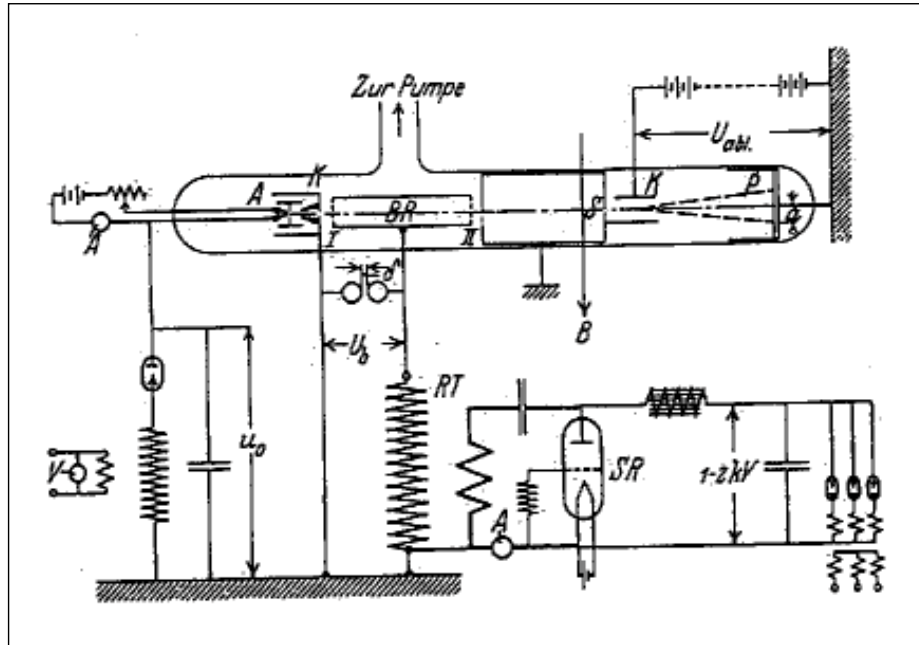
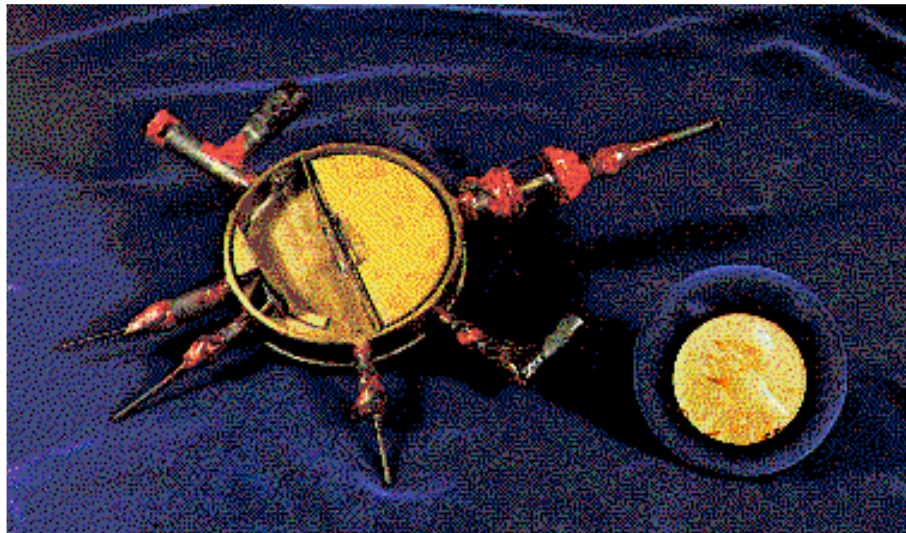Figure 1: Wideröe's single drift tube linear accelerator.

Figure 2: The first cyclotron with the Nobel Prize medal for comparison.

oscillators, and in 1949 delivered a 32 MeV 0.4 mA beam, 28 MeV of which resulted the linac sections requiring 2.5 MW of RF power. Soon thereafter, the Stanford Linear Accelerator Center embarked on its series of electron linacs culminating in the highly productive "two-mile" accelerator, initially at 20 GeV in the 1960s and later raised to 50 GeV. The 2856 MHz RF frequency adopted at SLAC was in the "S-band" of RADAR development.

1.2 The Synchrotron and Beam Dynamics

A way out of the cyclotron magnet problem was provided by the invention of the synchrotron by McMillan in the US and Veksler in Russia [4]. A time-varying magnetic field maintained a near-constant orbit radius as the beam energy was increased by repetitive passage through the RF accelerating field. The advances in RF technology associated with RADAR permitted higher frequency systems, and in the last half of the decade of the 1940s a veritable host of synchrotrons were constructed. My first meaningful accelerator experience was with the 300 MeV Cornell electron synchrotron. This device had a 1 m orbit radius, the guide magnetic field of which varied sinusoidally at 30 Hz with acceleration provided by a single 50 MHz RF cavity.

The "synchro" in synchrotron comes from the fortunate circumstance that the orbit radius remains (nearly) constant as the guide field and energy increase. This results from the *principle of phase stability*, and this is the first item in beam dynamics that we will look at in a little detail. Suppose acceleration is provided by a single resonator of negligible extent along the beam path at a frequency a multiple h, the *harmonic number* of the nominal orbit frequency. Then in the relativistic limit, the RF phase and particle energy at successive transits of the resonator may be written

$$\phi_{n+1} = \phi_n + \frac{2\pi h \alpha}{E_s} \Delta E_{n+1} \qquad (2)$$

$$\Delta E_{n+1} = \Delta E_n + eV(sin\phi_n - sin\phi_s) \qquad (3)$$

where eV is the maximum energy increment, the subscript s denotes a synchronous quantity (a value that pertains to a particle whose energy exactly matches the magnetic field), $\Delta E \equiv E - E_s$, and α is the ratio of a fractional change in orbit circumference to a fractional change in energy. The subscript n on the variables denotes values at the nth entry to the resonator.

Trajectories in $\Delta E/E, \phi$ phase space are plotted in Fig. 3 for typical values of the parameters. The rate of change of the magnetic guide field determines the rate of change of the synchronous energy, dE_s/dn, and this must be smaller than eV for a synchronous phase, ϕ_s, to exist. The frequency of these synchrotron oscillations is typically much less than the orbit frequency.

At the same time that the synchrotron principle was introduced, D. W. Kerst proposed the very elegant device that became known as the betatron, in which both steering and acceleration are provided by time-varying magnetic fields on and within the orbit. Though
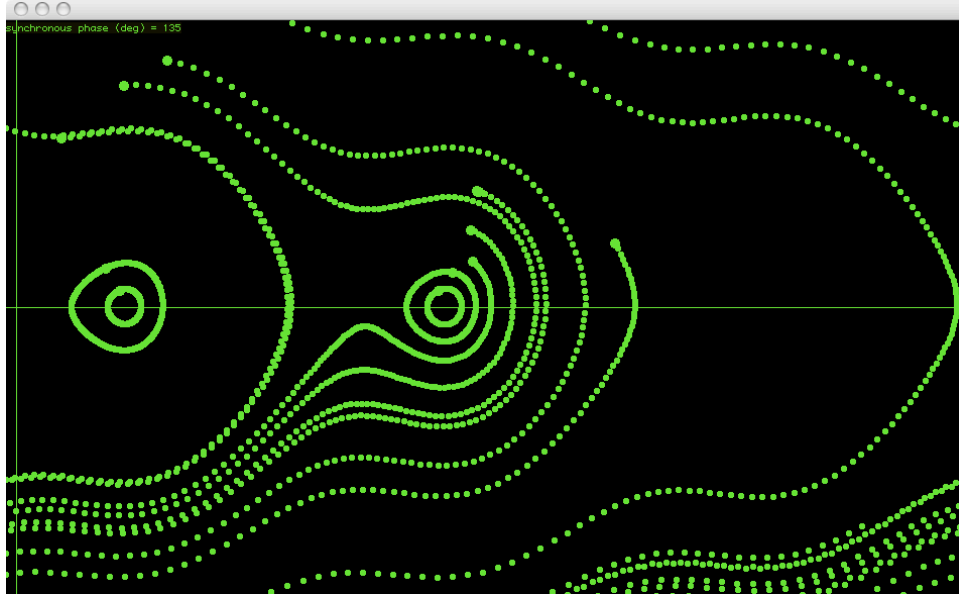
4

Figure 3: Trajectories in $\Delta E/E, \phi$ phase space.

electron accelerators of this design were constructed at the 300 MeV scale, the massive magnets precluded higher energies. But the main reason for mentioning this approach is that the transverse stability was treated in a famous paper by Kerst and Serber [5], and to this day these motions are referred to as betatron oscillations.

Unfortunately, even with a near-constant orbit radius, the aperture requirements of such high energy synchrotrons as the 3 GeV Cosmotron and Brookhaven and the 7 GeV Bevatron at Berkeley just scaled up with energy. This circumstance is illustrated in Fig. 4, in which an engineer is sitting at a desk within the vacuum chamber. The problem was a result of the weak focusing provided by the magnet systems. Think about a circular trajectory in the horizontal plane provided by a single magnet. In order to focus vertically the lines of field must bend outward. If they bend outward, the field as a function of radius must decrease, but it can decrease no more rapidly than the inverse of the radius, otherwise the necessary centripetal force for horizontal stability is lost. Therefore, for a particular angular divergence from the particle source, the amplitude of oscillation will be proportional to the radius of the device.

1.3 Alternating Gradient Focusing

In 1952, Ernest Courant, Stanley Livingston, and Hartland Snyder provided the way out of the aperture bind with their proposal of strong, or alternating gradient, focusing [6]. A quadrupole magnet provides focusing in one transverse direction and defocusing in the
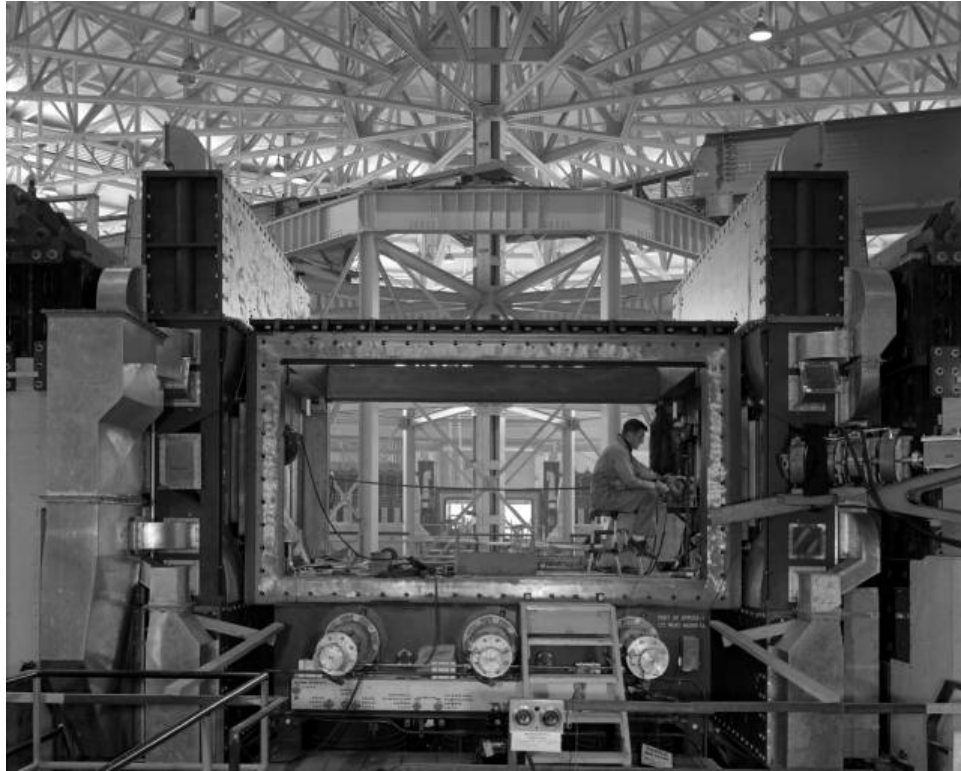
Figure 4: The Bevatron nearing completion.

other, but two quadrupoles of opposite focusing character can provide net focusing in both transverse senses. A sequence of alternating focusing and defocusing lenses can provide net focusing. If x and dx/ds represent the displacement and slope of a ray, then in the paraxial approximation, the progress of the ray through a pair of lenses of first focusing and then defocusing character, can be written in matrix form as

$$\left(\begin{array}{c} x \\ dx/ds \end{array} \right)_2 = M \left(\begin{array}{c} x \\ dx/ds \end{array} \right)_1 \tag{4}$$

where

$$M = \left(\begin{array}{cc} 1 & \ell \\ 0 & 1 \end{array} \right) \left(\begin{array}{cc} 1 & 0 \\ 1/F & 1 \end{array} \right) \left(\begin{array}{cc} 1 & \ell \\ 0 & 1 \end{array} \right) \left(\begin{array}{cc} 1 & 0 \\ -1/F & 1 \end{array} \right). \tag{5}$$

Here, F is the lens focal length and ℓ the lens separation. This separation matrix is a reasonable approximation to a dipole magnet for our purposes here. A similar matrix applies in the other transverse degree of freedom, with reversal of lens focusing character. Stability requires that M^n remain finite for arbitrarily large n, and an eigenvalue analysis quickly shows that this will be so provided $F < \ell/2$. For ℓ small compared with the circumference of the orbit, the wavelength of the transverse oscillation will be also small in comparison, and it is plausible though not yet demonstrated here that the amplitude of oscillation will also be reduced.

This step divorced aperture requirements from energy. Even before the Courant, Livingstone, and Snyder invention appeared in print, Robert Wilson at Cornell was determined to build an alternating gradient synchrotron. Construction of a 1.2 GeV electron accelerator of this design began immediately, and was in operation by the mid-1950s. The successes of the PS at CERN and the AGS at Brookhaven constructed later in the same decade are well known and they remain in operation today.

This was a time of good public support for particle physics. By the early 1960s enthusiasm and plans for a 200 GeV scale proton accelerator were well underway. and a US panel in 1963[7] recommended that this project receive highest priority. This resulted in the approval of Fermilab construction, and its first Director, Robert Wilson, immediately increased the design energy to an imaginative 500 GeV. A measure of his success is that routine performance at 400 GeV was routinely achieved. A proton synchrotron of similar scale and energy, the SPS, was constructed at CERN.

As time progressed and virtuosity in beam dynamics developed, emphasis for HEP turned toward colliders. The first electron-positron storage ring was ADA at Frascati[8], and for protons, this direction began with the ISR storage rings at CERN with their impressive 35 ampere beams[9]. Colliders will be discussed in more detail in the remaining talks. But a fitting acknowledgement to the progress of this period is the achievement of proton-antiproton collisions in the modified SPS at CERN, employing the technique of stochastic cooling[10]; this effort culminated in the discovery of the Z and W bosons in 1983.

2 Today

In order to continue to currently operating facilities, we should go into somewhat more detail regarding the beam dynamics[11]. But in contrast to the past, where accelerators drew for the most part on technology created for other purposes, in recent years substantial effort has been devoted to advancing technologies specifically for accelerator application and this development deserves comment.

All operating colliders are synchrotrons, and the following treatment reflects that circumstance. The event rate R in a collider is proportional to the interaction cross section σ_{int} and the factor of proportionality is called the *luminosity*:

$$R = \mathcal{L}\sigma_{\text{int}} \ . \tag{6}$$

If two bunches containing n_1 and n_2 particles collide with frequency f, then the luminosity is

$$\mathcal{L} = f\frac{n_1 n_2}{4\pi\sigma_x\sigma_y} \tag{7}$$

where σ_x and σ_y characterize the Gaussian transverse beam profiles in the horizontal (bend) and vertical directions and to simplify the expression it is assumed that the bunches are identical in transverse profile, that the profiles are independent of position along the bunch, and the particle distributions are not altered during collision. Whatever the distribution at the source, by the time the beam reaches high energy, the normal form is a good approximation thanks to the central limit theorem of probability and the diminished importance of space charge effects.

The beam size can be expressed in terms of two quantities, one termed the *transverse emittance*, ϵ, and the other, the *amplitude function*, β. The transverse emittance is a beam quality concept reflecting the process of bunch preparation, extending all the way back to the source for hadrons and, in the case of electrons, mostly dependent on synchrotron radiation. The amplitude function is a beam optics quantity and is determined by the accelerator magnet configuration. When expressed in terms of σ and β, a commonly used definition of the transverse emittance is

$$\epsilon = \sigma^2/\beta \ . \tag{8}$$

Of particular significance is the value of the amplitude function at the interaction point, β^*. Clearly one wants β^* to be as small as possible; how small depends on the capability of the hardware to make a near-focus at the interaction point. Equation 8 can now be recast in terms of emittances and amplitude functions as

$$\mathcal{L} = f\frac{n_1 n_2}{4\pi\sqrt{\epsilon_x \beta_x^* \, \epsilon_y \beta_y^*}} \ . \tag{9}$$

2.1 Beam Dynamics

A major concern of beam dynamics is stability: conservation of adequate beam properties over a sufficiently long time scale. Several time scales are involved, and the approximations used in writing the equations of motion reflect the time scale under consideration. For example, when we write the equations for transverse motion below no terms associated with phase stability or synchrotron radiation appear; the time scale associated with the last two processes is much longer than that demanded by the need for transverse stability.

Present-day high energy accelerators employ alternating gradient focusing provided by quadrupole magnetic fields. The equations of motion of a particle undergoing transverse oscillations with respect to the design trajectory are

$$x'' + K_x(s)\, x = 0 \ , \quad y'' + K_y(s)\, y = 0 \ , \tag{10}$$

with

$$x' \equiv dx/ds \ , \quad y' \equiv dy/ds \tag{11}$$

$$K_x \equiv B'/(B\rho) + \rho^{-2} \ , \quad K_y \equiv -B'/(B\rho) \tag{12}$$

$$B' \equiv \partial B_y/\partial x \ , \quad (B\rho) \equiv p/e. \tag{13}$$

The independent variable s is path length along the design trajectory. This motion is called a *betatron* oscillation because as noted above it was initially studied in the context of that type of accelerator. The functions K_x and K_y reflect the transverse focusing— primarily due to quadrupole fields except for the radius of curvature, ρ, term in K_x for a synchrotron—so each equation of motion resembles that for a harmonic oscillator but with spring constants that are a function of position.

These equations have the form of Hill's equations and so the solution in one plane may be written as

$$x(s) = A\sqrt{\beta(s)} \ \cos[\psi(s) + \delta], \tag{14}$$

where A and δ are constants of integration. In order that Eq. 14 be the general solution independent of δ, the phase and β satisfy

$$d\psi/ds = 1/\beta, \quad 2\beta\beta'' - \beta'^2 + 4K\beta^2 = 4. \tag{15}$$

The dimension of A is the square root of length, reflecting the fact that the oscillation amplitude is modulated by the square root of the amplitude function. In addition to describing the envelope of the oscillation, β also plays the role of an local $\lambda/2\pi$.

The wavelength of a betatron oscillation may be some tens of meters, and so typically values of the amplitude function are of the order of meters rather than on the order of the beam size. The beam optics arrangement generally has some periodicity and the amplitude function is chosen to reflect that periodicity. As noted above, a small value of the amplitude function is desired at the interaction point, and so the focussing optics is

tailored in its neighborhood to provide a suitable β^*. The second order equation for β simplifies considerably is differentiated once more to yield

$$\beta''' + 4K\beta' + 2K'\beta = 0. \tag{16}$$

In a drift space, the solution is a parabola resulting in the (eventually) familiar variation of β in the neighborhood of the interaction point.

The number of betatron oscillations per turn in a synchrotron is called the *tune* and is defined by

$$\nu = \frac{1}{2\pi} \oint \frac{ds}{\beta} \ . \tag{17}$$

Expressing the integration constant A in the solution above in terms of x, x' yields the *Courant-Snyder invariant*

$$A^2 = \gamma(s)\, x(s)^2 + 2\alpha(s)\, x(s)\, x'(s) + \beta(s)\, x'(s)^2 \tag{18}$$

where

$$\alpha \equiv -\beta'/2, \ \ \gamma \equiv \frac{1 + \alpha^2}{\beta} \ . \tag{19}$$

(The Courant-Snyder parameters α, β and γ employ three Greek letters which have many other meanings and the significance at hand must often be recognized from the context.) Because β is a function of position in the focussing structure, this ellipse changes orientation and aspect ratio from location to location but the area πA^2 remains the same.

As noted above the transverse emittance is a measure of the area in x, x' (or y, y') phase space occupied by an ensemble of particles. The definition used in Equation 8 is the area that encloses 15% of a Gaussian beam.

For present-day hadron synchrotrons, synchrotron radiation does not play a role in determining the transverse emittance. Rather the emittance during storage reflects the source properties and the abuse suffered by the particles throughout acceleration and storage. Nevertheless it is useful to argue as follows: Though x' and x can serve as canonically conjugate variables at constant energy this definition of the emittance would not be an adiabatic invariant when the energy changes during the acceleration cycle. However, $\gamma(v/c)x'$, where here γ is the Lorentz factor, is proportional to the transverse momentum and so qualifies as a variable conjugate to x. Therefore often one sees a normalized emittance defined according to

$$\epsilon_N = \gamma \frac{v}{c}\, \epsilon, \tag{20}$$

which is an approximate adiabatic invariant, e.g. during acceleration.

The motion described by Eq. 4 is a collection of straight line segments connected by angles at the quadrupoles, so one might wonder what is gained by introducing all this formalism. One advantage is that we can use familiar techniques associated with harmonic oscillators. Transform the variables in Eq. 14 according to $\zeta = x/\sqrt{\beta}$, $\phi = \psi(s)/\nu$, then ϕ

is an independent variable that increases by 2π in each turn, and the solution looks like an ordinary harmonic oscillator:

$$\zeta(\phi) = A cos(\nu\phi + \delta). \tag{21}$$

This maneuver was not invented for beam physics, by the way. It's called a Floquet transformation. Next suppose there is a difference, perhaps an error, in the magnetic fields used in the equations of motion Eq. 10. In the new coordinates, we have

$$\frac{d^2\zeta}{d\phi^2} + \nu^2\zeta = -\nu^2\beta^{3/2}\frac{\Delta B(\zeta, \phi)}{(B\rho)}, \tag{22}$$

where ΔB is the magnetic field perturbation, usually represented by a multipole expansion.

$$\Delta B = B_0(b_0 + b_1 x + b_2 x^2 + ...) \tag{23}$$

Here B_0 is some reference field strength; in a bending magnet, B_0 would be the nominal bend field. The coefficients b_n would represent dipole, quadrupole, sextupole and so on, field contributions for the various values of n.

With insertion of Eq. 23 into Eq. 22, the result is a driven harmonic oscillator with the consequence that resonance with some multipole may result if the tune is any rational number. That does not necessarily mean that oscillations will grow without bound for such tunes, for after all, the rational numbers are dense, but one ought to be careful. In a storage ring, one builds in some correction and adjustment magnets such as sextupoles and octopoles to control such behavior. For fixed target physics, nonlinear magnetic elements are invaluable for slow beam extraction.

2.2 Technology

The words "Armco A6 29 gauge punch type" were engraved in my head over a half-century ago because that was the steel one ordered if one wanted to build a 30 to 60 Hz magnet for a synchrotron. This designated a flat-rolled steel that the Armco Steel Corporation produced for transformers and other electrical equipment. It was certainly not developed with accelerators in mind.

As proton synchrotrons evolved into the several hundred GeV energy range, a magnet material capable of fields beyond the familiar 1.5 to 1.8 Tesla steel range became a very attractive goal. The magnetization, $\vec{M}$, of steel is limited, but superconductivity offered plenty of candidates for another route by high current. Although the discovery of superconductivity had been announced by Kammerlingh Onnes back in 1911, no industrial scale application had appeared. Here was an opportunity.

A magnetic guide field for an accelerator that did not rely on steel was not a completely original idea when the thought of superconducting material use became a high priority. Mark Oliphant, whose name is usually associated with magnetron development for RADAR, had initiated a 10 GeV proton synchrotron project in Australia in the early

1950s. The guide field was based on a model of two overlapping current-conducting circles, Each carries a uniform current density J though in opposite directions. If the centers of the two circles are separated by a distance d, the region of circle overlap has a uniform magnetic field of magnitude $\mu_0 J d/2$ and provides, due to absence of current, aperture for a beam pipe. The intended material was copper. This ambitious project was not completed due to the earlier achievement of comparable particle energies elsewhere.

But this picture of a bending magnet is interesting if we think of the current densities possible in superconducting materials. The immediate candidate was the NbTi alloy, a material with excellent mechanical properties for conductor development. The transition temperature is 9 K, and, with liquid He cooling, operation in the 4 K to 5 K region made sense. The potential J is up in the astronomical 3×10^9 ampere/m^2 neighborhood. If we think we can make a cable capable of 10% of that current density, then two circles with centers separated by about 2 cm may provide a field in the overlap region of about 4 T.

Scarcely had the 400 GeV project at Fermilab been completed than Robert Wilson announced his intention to design a 1 TeV accelerator based on superconducting bending magnets to occupy the same enclosure as its predecessor. This device, initially called the Energy Doubler and ultimately known as the Tevatron, was to occupy the same enclosure as its predecessor but achieve 1 TeV. The bend field needed would be 4.4 T. Magnet development began in 1972, with field derived directly from conductors in a manner rather like the configuration of two-circles mentioned above.

In a conventional iron-based magnet, the field is shaped by machining or lamination stamping, operations capable of high precision. The achievement of corresponding tolerances in conductor placement is a significant challenge. But at least in the case of copper conductors, useful examples can be found in just about any hardware store. Not so with NbTi; literally years of effort went into the developing superconducting cable for the magnets.

The problem was that the material will switch from the superconducting to the normal state due to an energy deposition of only about 1 mJ/gm as a result of the very low specific heat of this (and other) materials at low temperature. This transition is called a "quench". Energy sources at this level are easy to find. Particle beams with energies of more than a megajoule are obvious candidates, but conductor movement at the 1μm level in a 4 T field will also provide such heating. Just about the only substance with a useful heat capacity in the 2-4 K temperature range is liquid Helium, and so pervasive association with this coolant is necessary, as is the close proximity of some conductor, such as Cu, that maintains low resistivity in its normal nonsuperconducting state. It was an impressive achievement that a cable based on NbTi was developed that provided about 10% of the current density of the pure superconductor.

The other aspect of the superconducting accelerator that would come to mind – the refrigeration suitable for the 4 K environment and its liquid helium coolant flow – was already a mature technology thanks to the demand of helium for heliarc welding. I was surprised to learn that the gas was liquefied near the natural gas wells in the US Southwest
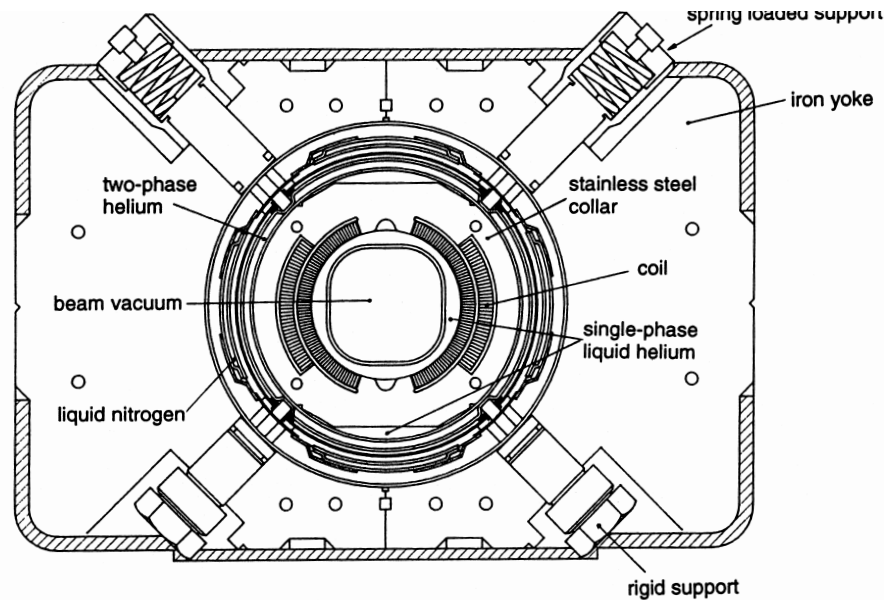
Figure 5: Tevatron dipole magnet.

and shipped in sealed dewars for use elsewhere in the world.

The Tevatron was designed as a fixed target accelerator, and so a reasonably rapid cycle time was important. A 60 s cycle time at 1 TeV was set as the goal with the intent that the event rate in the neutrino program would be about the same as that familiar from the preceding 400 GeV synchrotron with its 10 s cycle. The magnet cross section is shown in Fig. 5. A variety of cooling, eddy current, and space occupancy concerns resulted in a small magnet cross section with room-temperature iron surrounding the coil cryostat. This design was not without compromises, but that is a long story that would take up too much space here.

Magnets for storage rings need not be designed with corresponding concern for rate-of-change of field. While the Tevatron magnet was undergoing development, DESY launched the HERA (Hadron-Electron-Ring-Accelerator) project to collide 1 TeV protons with 30 GeV electrons using two synchrotrons, one for each particle species, in the same underground tunnel. The cross section of the superconducting bending magnet for the HERA proton ring is depicted in Fig. 6. The cryostat encompasses the entire magnet material. The magnet is capable of excitation somewhat above the field corresponding to 1 TeV particles.

The other aspect of superconducting technology that was already receiving great attention during this time was its application to RF systems. I would like to defer discussion on this topic till the next two lectures.

To conclude this lecture, it might be useful to contrast the early accelerators of Fig. 1.1
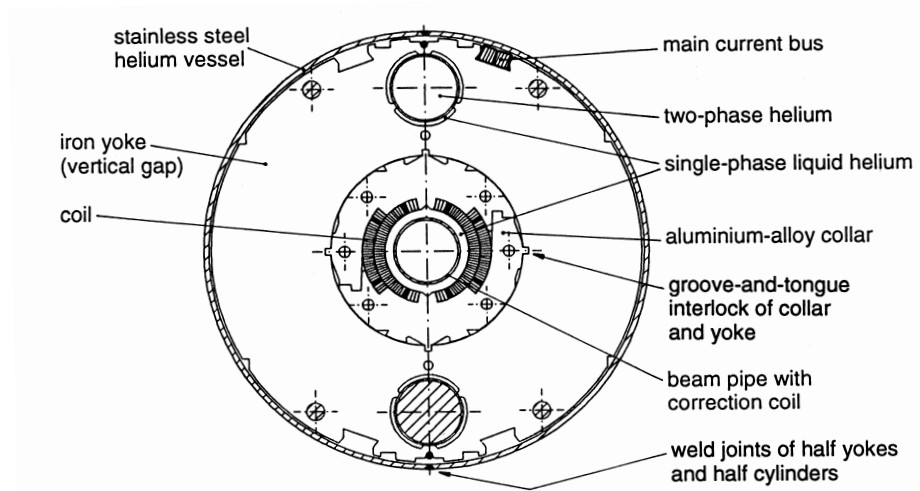
Figure 6: HERA proton ring dipole magnet.

and Fig. 1.1 with the accelerator complexes of Fermilab and DESY. A view of Fermilab is shown in Fig. 7 with its collection of accelerators necessary to $p\bar{p}$ collisions and fixed target neutrino experiments. The DESY layout appears in Fig. 8. In order to provide electron-proton collisions, DESY must build two of just about everything.

3 Tomorrow

The Large Hadron Collider (LHC) dominates the horizon. This proton-proton collider under construction at CERN promises a cms energy of 14 TeV and a luminosity two orders of magnitude above that reached at the Tevatron. Like the Tevatron, there was an accelerator enclosure constructed for an earlier project, the very successful LEP electron-positron storage ring. This 27 km tunnel is just over a factor of four greater in circumference than that occupied by the Tevatron, so just with replication of the superconducting magnet technology developed at Fermilab and DESY a significantly higher energy would be possible. But in the earliest design discussions, a bend field about twice as large – in the 8 T range – was adopted, implying the seven-fold or so increase in collision energy.

The CERN accelerator complex is shown in Fig. 9. Protons destined for injection into the LHC begin their history in a hydrogen plasma, from which they are extracted by electrodes applying a potential difference in the 50 kV range. They are then accelerated to 750 keV by a device called a radiofrequency quadrupole (RFQ), which is today's replacement for the Cockcroft -Walton electrostatic device, to get particles up to a speed suitable for
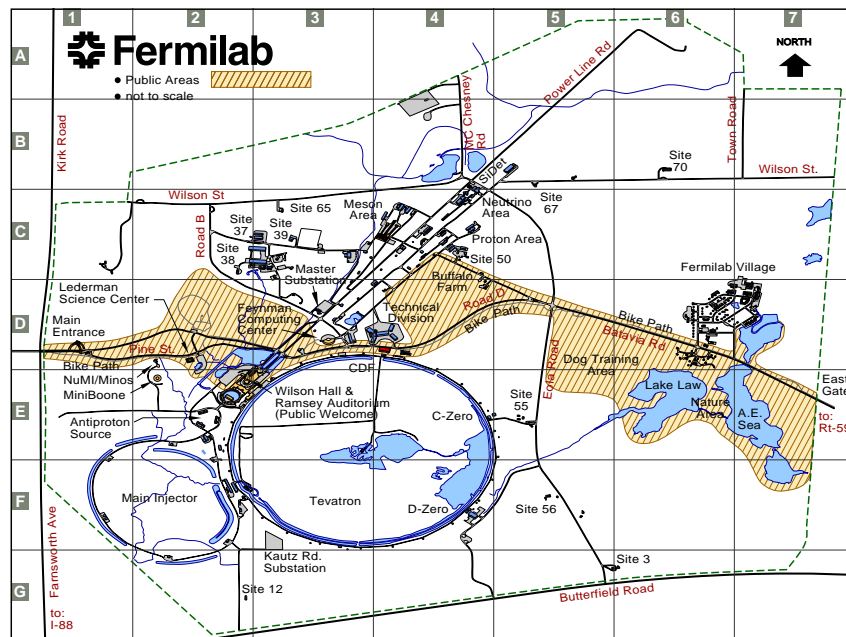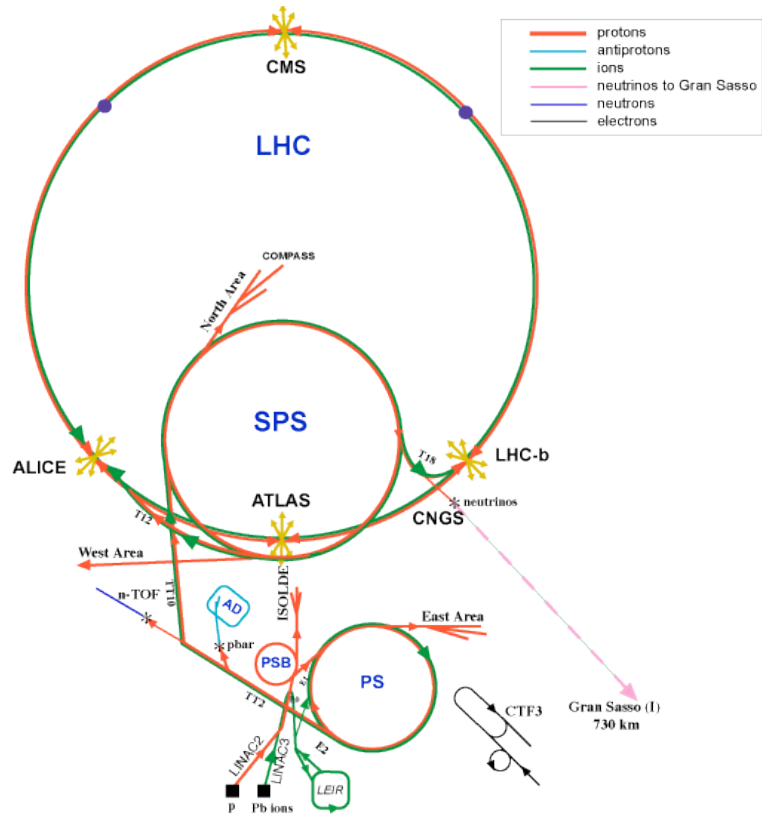
Figure 7: Fermilab site.

Figure 8: DESY site.

Figure 9: The CERN accelerator complex.

acceptance by a linear accelerator, named LINAC-2. This linac delivers 50 MeV protons to the Proton Synchrotron Booster (PSB in the diagram) which accelerates the particles to 800 MeV for acceptance into the famous and venerable Proton Synchrotron (PS), the 28 GeV protons from which have served the CERN program since about 1958.

The next-to-the-last step in the journey is the Super Proton Synchrotron (SPS), which, if I remember correctly, is the last of the big proton rings constructed using conventional steel and copper magnets. The role of the SPS as its SP$\bar{\mathrm{P}}$S incarnation in the W and Z boson discoveries has been noted earlier. The versatile SPS served as an electron accelerator in the LEP years, and its next mission is the delivery of 450 GeV protons to LHC.

It is tempting to go into a story about why this chain of accelerators is necessary, at CERN or anywhere else. But much of it is history of a particular laboratory, and the design trade-offs that have been made in the past. A few rules are clear, such as one should not attempt injection into a synchrotron at a magnetic field in the neighborhood of 10 gauss. The dominance of remanence makes any conception of field quality irrelevant. At a more sophisticated level, the emittance handoff and implications for aperture from stage to stage is important.

For this talk, I would like to concentrate on the LHC ring itself, and see if some of the material presented in the preceding three talks can be helpful in discussing the design. I will remark first on the parameters and then comment on technology. Extensive material on this subject may be found on the LHC website at CERN[12].

3.1 LHC Design Parameters

For proton-proton collisions, two magnet rings are needed, and the bunches collide with a crossing angle. So some modification of Eq. 9 is needed, and for the LHC the luminosity expression takes the form

$$\mathcal{L} = f \frac{n^2}{4\pi\epsilon\beta^*} F \tag{24}$$

where now the bunches for the two particle species are assumed to be equally populated, the transverse emittances are the same as are the amplitude functions in the transverse degrees-of-freedom. The factor F reflects the reduction due to the crossing angle θ_c:

$$F = \frac{1}{\sqrt{1 + \left(\frac{\theta_c \sigma_z}{2\sigma^*}\right)^2}} \tag{25}$$

where σ_z is the rms bunch length and σ^* the transverse rms beam size at the interaction point. The influence of F on the luminosity goal is not particularly dramatic; the factor is about 0.84, necessary to mention for completeness but not a big impediment with regard peak luminosity.

The emittance, ϵ, is a major consideration in determining beam path aperture requirements, as well as occupying a place in the denominator of Eq. 24, therefore there is strong

motivation to accept the smallest value that the injection system can deliver. For the CERN complex, the normalized emittance, $\epsilon_n = \gamma\epsilon$, is about 4×10^{-6} m.

Generally speaking, it is desirable to achieve the design luminosity with as few particles as necessary. The number of protons per bunch is limited by the beam-beam tune shift. The particles of one bunch present a nonlinear lens to particles of a bunch with which it collides, which leads to a tune spread given by

$$\xi = \frac{nr_p}{4\pi\epsilon_n}. \tag{26}$$

Here, $r_p = 1.5 \times 10^{-18}$ m is the classical radius of the proton. There is an extensive literature on how large the tune-shift parameter, ξ, can be. For one thing, it means there is a spread in the batatron oscillation tunes, hence the possibility of overlap with resonances best avoided. For proton-proton collisions, a limit of about 0.015 is considered reasonable, and since there were to be several interaction regions in the LHC, the designers chose a value of 0.005 for each bunch-bunch passage. This choice leads to $n = 10^{11}$ protons per bunch.

The attempt to achieve a near-focus at the interaction point is reflected in the choice of β^*. The smaller β^* is, the larger the angular divergence at the interaction point, and therefore the larger the beam becomes elsewhere. For the high luminosity regions the designers chose $\beta^* \approx 0.5$ m as consistent with optics design and detector space requirements.

The final parameter that remains in Eq. 24 is the collision frequency. In order to yield the target luminosity, using the numbers above, we get a collision frequency of $\pi \times 10^7$ Hz. The 27 km circumference gives a 11 kHz rotation frequency for relativistic particles. The ratio yields the number of bunches, about 2800. Note that although 10^{34} cm^{-2}s^{-1} is two orders of magnitude above what has been achieved in earlier hadron-hadron colliders, such is the progress over the years that no single parameter strains credibility.

If I had to choose one consequence of the basic parameters that concerns me most, I would select the stored energy per beam. Use of CERN's more exact figures, one arrives at 362 MJ for each proton beam. For comparison, a kilogram of TNT delivers about 4 MJ to an explosion. At the other end of the spectrum of magnitudes is the few mJ/g quench sensitivity of the conductor. The LHC is scheduled for initial operation at reduced luminosity in order to deal with the beam power and other inevitable start-up problems.

The space constraints of the CERN tunnel do not permit two separate cryostats for the superconducting magnets for the two beams. A cross-section of the LHC bending magnet is characterized in Fig. 10, which illustrates the magnetic field lines. The fields in the two beam paths are inescapably linked, unlike the situation in earlier accelerators. The white circles represent holes in the steel necessary to field accuracy, and reflect the detailed design software needed for this effort.

Most of the path length in the tunnel will just be alternating gradient structures, repetitive cells of quadrupoles and bending magnets such as introduced in Eq. 4. Over the years since 1952 when the advent of alternating gradient focusing broke the aperture-
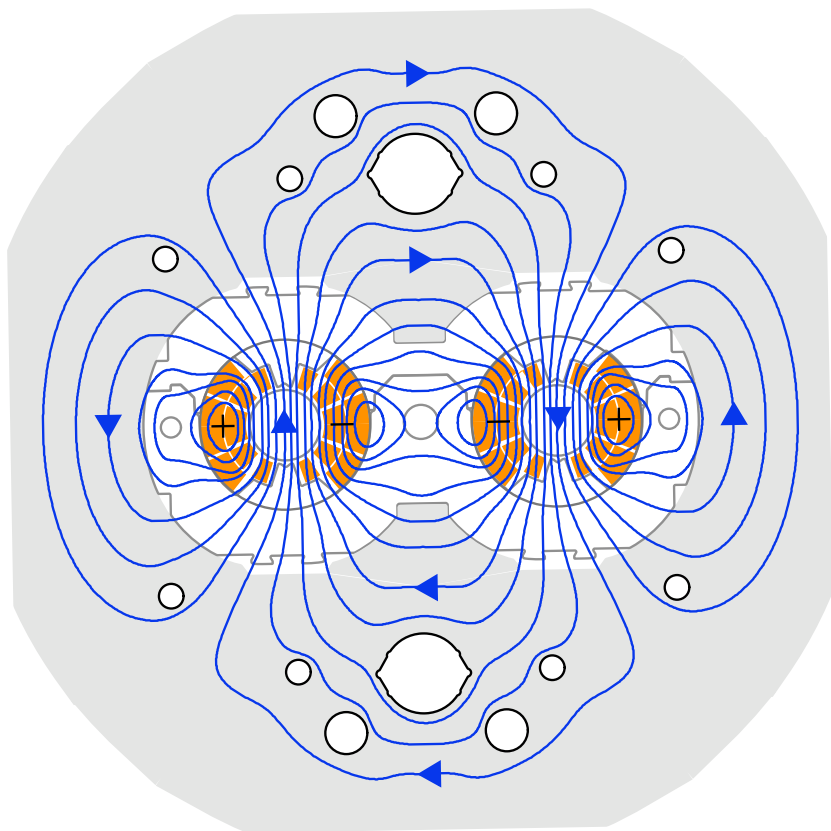
Figure 10: The two aperture bending magnet for the LHC.

proportional-to-circumference requirement, beam pipes have still grown in diameter but less rapidly, more like the square root of the circumference. I don't know how the discussions went concerning the length of the alternating gradient cells for the LHC, but I suspect that the combination of the space constraints of an existing facility and an appropriately more conservative approach to aperture usage lead to the adoption of an alternating gradient half-cell length of 53.5 m, less than twice the 30 m of the Tevatron although a circumference ratio of over 4 could have been used as an argument for a different choice. I mention this aspect of the subject only to suggest that there is flexibility in design choice, and the interplay among design choices, cost optimization, and performance is an important part of the process.

3.2 LHC Technology

The superconducting magnets of the LHC are based on NbTi technology. It's interesting that despite the very attractive parameters of Nb_3Sn such as an 18 K transition temperature, this difficult brittle material has yet to find broad application. It was again promoted for use at LHC as it was decades ago for the Tevatron. With use of NbTi, some aspects of the LHC magnets such as the cable are relatively familiar. Two significant differences are the necessity of 2 K cryogenics and the mechanical stresses resulting from the 8.33 T magnetic fields. Recalling that a sheet current with 0.5 T on one side and 0 T on the other sustain a 1 atmosphere pressure difference, the 8.33 T implies a hoop stress of about 275 atmospheres. The excellent field quality of the LHC bending magnets attests to the success in their design and construction.

The Tevatron was commissioned as a collider over 20 years ago, and has undergone continual improvement. It is quite likely that the LHC will last considerably longer as a productive research instrument, and that it will benefit from advances in superconducting technology already underway, such as the Nb_3 quadrupoles mentioned in Sec. 4.3.1 below. Luminosity at the 10^{34} cm^{-2}s^{-1} is almost certainly not an upper bound, yet it may be expected that progress in computer technology will keep pace with the growth in data rate.

4 The Future

Present emphasis for a next major HEP facility is on an electron-positron linear collider to complement the LHC with its hadron basis. Earlier I have not said much about electron accelerators because their particular problem, synchrotron radiation, did not intrude upon the hadron world, except as noted above in the context of the LHC vacuum chamber design. But at this point I should talk a bit below on synchrotron radiation. Design, approval, and construction of an International Linear Collider is a near term priority, and there are some comments in a section under that title. Finally, there is the matter of the longer term and that is the subject of the concluding section.

4.1 Synchrotron Radiation

Suppose an electron is traveling on a circular path. With respect to an inertial frame moving tangent to the circle at the same speed as the electron, the electron will be briefly at rest at the point of fangency. So in the inertial frame, we can apply the Lorentz formula for the radiated power:

$$P' = \frac{1}{6\pi\epsilon_0}\frac{e^2 a'^2}{c^3} \tag{27}$$

where the primes denote quantities measured in the speeding inertial frame. In the laboratory frame, $P = P'$ and $a = \gamma^2 a'$. For a relativistic electron, $a = c^2/\rho$, where ρ is the radius of curvature. The radiated power becomes

$$P = \frac{1}{6\pi\epsilon_0}\frac{e^2 c}{\rho^2}\gamma^4. \tag{28}$$

In a synchrotron that has a constant radius of curvature, the energy lost due to synchrotron radiation is the above expression multiplied by the time spent in bending magnets, $2\pi\rho/c$. In familiar units, we then have

$$W = 8.85 \times 10^{-5} E^4/\rho \;\; \text{MeV per turn} \tag{29}$$

where the energy E is in GeV, the radius ρ is in km, and the expression is evaluated for electrons. The correspondence of the "8.85" here with the "8.85" in ϵ_0 is interesting but accidental; they differ in the next digit.

For some years, this fourth power law and its consequences was regarded as an aggravation by electron synchrotron designers, and it does indeed set an effective limit to the energy of this variety of accelerator. For LEP at 100 GeV the energy loss per turn was xxx MeV, which required compensation by the RF acceleration system. At this point, a linear collider becomes an attractive alternative, and this recognition lead to the SLAC Linear Collider, and ultimately to the present interest in an International Linear Collider as the lepton counterpart to the LHC.

That synchrotron radiation was not all bad was also recognized quite quickly. As a graduate student, I assisted Hartman and Tomboulian[13] in 1952 on their study of radiation from the Cornell 300 MeV synchrotron as a potential source useful to solid state physics. Their experiment was a precursor of the explosion of interest in synchrotron radiation that continues today, with linear accelerator based free electron lasers as prime examples. Synchrotron radiation also benefits today's e^+e^- circular colliders such as PEP II and KEKB as the damping process permits continuous injection, "topping up", to maintain the luminosity in the stored beam.

The use of damping rings to provide emittance reduction and high transverse emittance ratios was anticipated by the important work of K. Robinson.[14], in which he arrives at the equilibrium distributions resulting from the balance between radiation damping and oscillation excitation, the latter arising from quantum fluctuations in the radiation rate.

In the proton synchrotron world, synchrotron radiation could be ignored until the energy scale of the LHC. For protons, the expression corresponding to Eq. 29 is

$$W = 7.8 \times 10^{-3} E^4 / \rho \ \text{keV per turn} \tag{30}$$

where E is now in TeV and ρ remains in km. For the LHC, synchrotron radiation presents a significant load to the cryogenic system, and impacts magnet design due to gas desorption and secondary electron emission from the wall of the cold beam tube.

4.2   International Linear Collider (ILC)

There had been a general understanding in the HEP community for some years that the natural complement to the LHC would be an electron-positron linear collider at a cms energy of 0.5 to 1 TeV. Two technologies — "warm"' and "cold" — competed for acceptance, each with vocal and persuasive adherents. The copper-based RF structures so successfully used for decades, particularly in the highly successful SLAC program, was a natural contender. The "cold" approach using superconducting cavities had developed rapidly in recent years and had emerged as the alternative. In August 2004, a distinguished international panel of unbiased physicists opted in favor of the "cold" approach. This was certainly not an easy decision, and the arguments one way or another make interesting and valuable reading.[15]

With the technology choice in place, significant progress has been made in pulling together past design efforts into form suitable for world-wide collaboration. The Global Design Effort (GDE) has a website[16] on which developments are posted and the goal is to present a design report by the end of 2006. A useful file to scan — a recent version amounts to 305 pages — is the Baseline Configuration Document which is to be found on the GDE website.

The luminosity choice at 2×10^{34} cm^{-2}s^{-1} is similar to that of the LHC. The energy is lower; the plan has two 250 GeV linacs colliding initially with attention given to expansion to twice that energy. Yesterday, we went through the luminosity calculation for the LHC. It is interesting to scale from the parameters of the LHC to the present circumstance. Let me rewrite Eq. 9 in terms of the normalized emittances.

$$\mathcal{L} = \gamma f \frac{n_1 n_2}{4\pi \sqrt{\epsilon_{x,n} \beta_x^* \, \epsilon_{y,n} \beta_y^*}}, \tag{31}$$

and below we will assume that the bunch populations are the same: $n_1 = n_2 = n$. Unlike a synchrotron in which the bunches return to the interaction point after every circuit of the ring, in the linac the energy used to accelerate the particles is lost with each passage and must be reinvested anew. The 362 MJ in the LHC orbit becomes 12 MJ at ILC energy, and that gets multiplied by the 10^4 Hz orbit frequency to give a power requirement of 120×10^3 MW, or the output of 120 large power plants. That is obviously absurd, so the product fn must be reduced by 3 orders of magnitude.

23

Beam-beam effects argue for a reduction in the bunch charge by about an order of magnitude. So try a reduction in the collision frequency by a factor of 100, to bring the power in bounds. In the luminosity, this last 100 is largely compensated by the larger γ of the leptons. But from n^2 we are still down by 100, and this has to be made up from the emittances and amplitude functions. From Table 1.1 of the 16 March 2006 version of the Basic Configuration Document, $\epsilon_{y,n} = 0.04$ mm-mrad, $\beta_x = 21$ mm, and $\beta_y = 0.4$ mm for one of the "nominal" configurations. From elsewhere in the document, I infer $\epsilon_{x,n} = 4$ mm-mrad. These are parameters of the ILC that strike me as adventurous, but fortunately major progress down this road had been made at the SLAC Linear Collider, the first of its species. From the collider parameter tables in the 1998 Review of Particle Physics, we see for the SLC $\epsilon_{y,n} = 5$ mm-mrad, $\beta_x = 2.5$ mm, and $\beta_y = 1.5$ mm.

The two beam-beam effects that should be mentioned are (the unfortunately bilingual) beamstrahlung and disruption. The first of these refers to the electromagnetic radiation from particles of one bunch due to the fields in the other. The procedure that was folllowed to obtain Eq. 29 may be used here, at least as a beginning. That is, look in the rest frame of a particle that does the radiating as a bunch charges by in the opposite direction. That approach gets some of the dependencies right, but the averaging over distributions is too complicated for me. The result for the average fractional energy loss is[17]

$$\delta_E \approx 0.86 \frac{r_e^3 n^2 \gamma}{\sigma^z (\sigma_x^* + \sigma_y^*)^2} \tag{32}$$

where $r_e = e^2/(4\pi\epsilon_0 mc^2)$ is the classical radius of the electron. An ambiguity in the collision energy is not what one wants — the lack of such uncertainty is what is hoped for in an e^+e^- collider — so δ_E is limited to a few percent in the design. By making $\sigma_x \gg \sigma_y$, one controls δ_E and adjusts the luminosity by σ_y.

Disruption refers to the pinch effect that occurs when two oppositely directed bunches of the opposite sign of charge overlap and the magnetic field acts to compress the bunch density and hence to enhance the luminosity. The disruption parameter, D, is the ratio of σ_z to the focal length of the compression lens, and for the ILC design $D \approx 25$ which brings to mind an interesting picture of particle passage through the overlap region. The associated luminosity enhancement factor, H_D, contributes about a factor of two.

Finally there is the so-called hourglass effect, characterizing the circumstance that the beam size grows rapidly with longitudinal distance form the interaction point. The amplitude function in the neighborhood of the IP is a parabola: $\beta(z) = \beta^* + z^2/\beta^*$, therefore there is very limited gain in reducing β_y^* below σ^z. Combination of the above statements about beam power, beam-beam concerns, and the hourglass effect yields the interesting luminosity expression for a linear collider found in the TESLA Design Report[17]:

$$\mathcal{L} \approx 5.74 \cdot 10^{20} \ m^{-3/2} \ \times \ \frac{P_b}{E_{cm}} \times \left(\frac{\delta_E}{\epsilon_{y,n}} \right)^{1/2} \times H_D \tag{33}$$
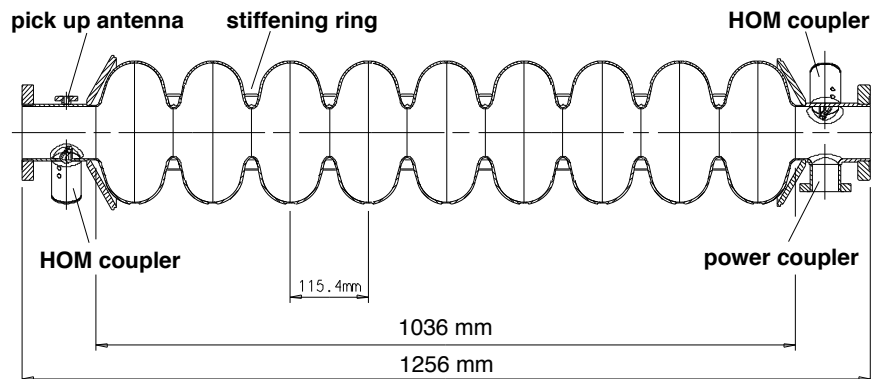
Figure 11: The 1.3 GHz TESLA cavity.

The accelerating structures in the current ILC model are the result of the very successful TESLA Collaboration technology development. In 1992, Bjørn Wiik started this effort, initially centered at DESY, with the target of an increase of the accelerating gradient to 25 MV/m from the 5 MV/m or so available from superconducting cavities at that time. A measure of the progress may be found with reference again to the Baseline Configuration Document where the current design gradient is given as 31.5 MV/m based on demonstrated performance above this level.

The TESLA structure is shown in Fig. 4.2. The shape differs from that of the disk-loaded waveguide in order to increase the ratio of the accelerating field on the axis to the magnetic field to that on the material surface. Niobium is a Type I superconductor, hence will revert to the normal state ("quench") if the magnetic field at the surface exceeds the limit, for Nb, of about 0.2 T. The fabrication techniques are familiar from sheet-metal technology, for Nb is readily worked by standard machining methods. A particular aspect requiring special attention is preparation of the RF surface to eliminate particles or protrusions as sites for field emission. Improvements are still being made as this is written. The ready availability of the "clean-room"' environment developed by the pharmaceutical and microchip industries has been of great value. I commented in an earlier lecture about the convenience of finding conductors and steel almost in the nearby hardware store; now clean-rooms are cataloq items, ultra-pure water systems can be rented from you local water source, and a 100 atmosphere rinsing pump is available at any supplier of car wash equipment. The development of surface preparation techniques in the pursuit of higher gradient continues, as does the investigation of other resonator shapes.

4.3   R&D in Particle Accelerators

Three directions come immediately to mind: improvement of existing accelerators; new accelerators such as muon synchrotrons based on present principles; and to borrow an expression from Monty Python "now for something completely different". I will just comment on a few examples, but to get a broad perspective, take a look at the annual report prepared by the Department of Energy Office of High Energy Physics which contains summaries of the many activities supported by that agency[18].

4.3.1   Improvements

Back when the Tevatron program was in its infancy about 1972, there were urgings to consider Nb_3Sn for the magnet coils. A premature suggestion then, but now that the technology of this material has advanced considerably, some limited application in LHC has been made, and an improvement program, LARP[19] is developing high gradient quadrupoles based on Nb_3Sn for a luminosity upgrade. The 18 K transition temperature with its implications for higher magnetic fields makes this material attractive, but it would be nice to find a substance with a transition temperature above the boiling point of liquid hydrogen; a possible candidate is MgB_2, which shows potential for technological development as a conductor. To date, application of the the really high T_c, above liquid nitrogen, materials has not been possible.

The small emittances needed for the ILC are not achieved easily, and the associated damping rings have significant cost and operational impact. A low-emittance source of polarized electrons is an interesting direction of study, and here the photoinjector RF electron gun is a candidate. Whether or not the combination of normalized emittance at the level of the ILC requirements and polarization in the confined vacuum environment of the RF electron gun is yet to be determined. That a high transverse emittance ratio can be provided by an RF electron gun has already been demonstrated[20]. Polarization is a different matter; the GaAs cathodes needed for polarization require an excellent vacuum environment at the 10^{-12} Torr level that is not readily obtained in an RF gun. Incidentally just as another example of technology obtained from elsewhere, GaAs is the cathode material used in night-vision goggles.

4.3.2   Muon Rings

The synchrotron radiation formula Eq. 29 was called the "Law of Inclemency" many years ago by my colleague John Rees regarding the future of electron synchrotrons. Muons present themselves as alternatives to electrons; they are still leptons without the burden of substructure, the mass is two orders of magnitude above that of the electron with great benefit to the γ^4 term, and perhaps their limited lifetime in their rest frame of 1.6 μs will not prove to be an impediment to their usage provided high enough magnetic guide fields are available. Actually, in calculating the number of turns that a muon will survive in a
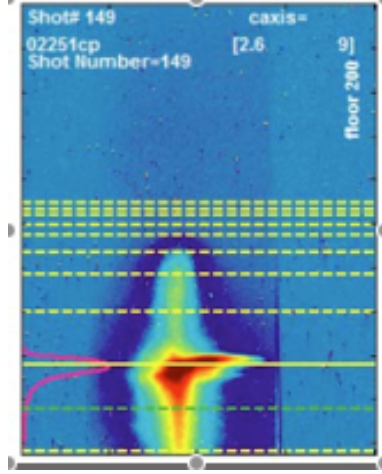
Figure 12: Doubling of 28.5 GeV in plasma wakefield. The solid yellow line is at 30 GeV, the dotted lines above are at 10 GeV intervals.

ring, the γs cancel out with the result $n \approx 300B$, where B is the average value of the guide field.

There is an extensive literature on muon colliders and storage rings[21], and I will not attempt to discuss these complex designs here. The subject has attracted bright inventive people, and their thoughts have been beneficial to other aspects of accelerator physics. For example, an experiment currently in development for the photoinjector at Fermilab is an atttempt to produce very low emittance electron bunches directly from the source without reliance on a damping ring. There is a significant international R&D underway on muon accelerator issues, concentrating as to be expected on emittance reduction following the large phase-space occupancy at the muon production source. Given the muon lifetiime, methods such as electron and stochastic cooling are not candidates and so ionization cooling has been given prominence.[22]

4.3.3 A New Direction?

Just two weeks or so ago, we learned that the E167 experimenters[23] using particles at 28 GeV from the SLAC linac had observed doubling of the energy of some of the bunch as a result of the plasma wakefield process in a 90 cm plasma column. This result is shown in Fig. 12. There have been great hopes attached to plasma research for more than the half-century since the initiation of the controlled fusion programs of the 1950s. Now not only is the power generation direction progressing as exemplified by ITER, but now we have a splendid example of acceleration within a plasma.

In the plasma wakefield process, an electron bunch initiates a strong plasma wave that can accelerate particles behind the head of the bunch or particles in a trailing bunch. An

27

order of magnitude estimate of the energy gain may be made as follows. Represent the electric field in the wave by

$$E(x, t) = E_0 cos[k(x - vt)] \tag{34}$$

where x is in the direction of motion of the bunch. From Maxwell's equation $divE = \rho/\epsilon_0$, we have

$$E_0 = \frac{ne}{\epsilon_0 k} = \frac{nec}{\epsilon_0 \omega_p} \tag{35}$$

where n is the electron density in the plasma, ω_p is the plasma frequency, and c (over)estimates the phase velocity of the wave. For $n = 10^{23}/\text{m}^3$, the result is about 30 GeV/m.

This is a big step, but it does not solve all of our problems just yet. We still have Eq. 33 to contend with. There is still the power problem. But why not recapture the beam power after collision? That is the mission of the Energy Recovery Linac investigations[24], and then the power in the equation may just represent the inefficiency of the recovery process. How emittance can be preserved within the plasma environment is another interesting question.

But the point I wish to make is that there is evidence of a new departure. It is much too early to guess how this may play out. This plasma related physics is alluring, and may point a way to relieve the cost-and-scale difficulties that hinder further progress in the familiar directions of the past.

References

[1] "Life and Work of Rolf Wideröe", Compiled and Edited by Pedro Waloschek, downloaded January 2006 from http://www.waloschek.de/pedro-texte/wid.htm.

[2] Ernest O. Lawrence and M. Stanley Livingston, Phys. Rev. **38**, 834 (1931).

[3] http://lawrencehallofscience.org

[4] E. M. McMillan, Phys. Rev. **68**, 143 (1945). V. I. Veksler. Comptes Rendus de l'Academie Sciences de l'URSS **43**, 8 (1944).

[5] D. W. Kerst, R. Serber, Phys. Rev. **60**, 53 (1941).

[6] E. D. Courant, M. S. Livingston, H. S. Snyder, Phys. Rev. **88**, 1190 (1952).

[7] The history of the deliberations surrounding the formation of the then-named National Accelerator Laboratory may be found at http://history.fnal.gov/GoldenBooks/gb_hoddeson.html.

[8] http://inf.infn.it/accelerator/ada/;

[9] www.pubmedcentral.nih.gov/ articlerender.fcgi?artid=433316

[10] From the original period, see for example D. Möhl *et al*, Physics Reports 58, No. 2 (1980).

[11] There are many texts on the subject, including one of which I am a coauthor, but it is quicker to go got to http://cas.web.cern.ch/cas/ which is the home-page of the CERN Accelerator School, where there are links to other sources of information, where all of this material and more can be found.

[12] http://lhc.web.cern.ch/lhc/

[13] P. L.Hartman, D. H. Tomboulian, Phys. Rev. **87**, 233(1952)

[14] K. W. Robinson, "Radiation Effects in Circular Electron Accelerators", Phys. Rev. **111**, No. 2 (1958).

[15] U. S. Linear Collider Technology Options Study, March 1, 2004.

[16] http://www.linearcollider.org/

[17] http://tesla.desy.de/tdr/

[18] Advanced Accelerator Research Annual Report 2005, DOE Office of High Energy Physics

[19] http://www.agsrhichhome.bnl.gov/LARP/

[20] D. Edwards et al., in Proceedings of the XX International Linac Conference, Monterey, CA, 2000 (SLAC, Stanford, 2000), pp. 122124.

[21] See for example http://www.cap.bnl.gov/mumu/studyii/

[22] prst-ab.aps.org/abstract/PRSTAB/v8/i6/e06100

[23] Material on this series of experiments may be found at http://www.slac.stanford.edu/grp/arb/siemann.pdf

[24] http://erl.chess.cornell.edu/